arXiv:1403.7791v1 [cs.DC] 30 Mar 2014

POSH: Paris OpenSHMEM
A High-Performance OpenSHMEM Implementation for Shared Memory Systems

Camille Coti *
LIPN, CNRS-UMR7030,
Université Paris 13, F-93430 Villetaneuse, France
camille.coti@lipn.univ-paris13.fr

Abstract

In this paper we present the design and implementation of POSH, an Open-Source implementation of the OpenSHMEM standard. We present a model for its communications, and prove some properties on the memory model defined in the OpenSHMEM specification. We present some performance measurements of the communication library featured by POSH and compare them with an existing one-sided communication library. POSH can be downloaded from http://www.lipn.fr/~coti/POSH.

1 Introduction

The drive toward many-core architectures has been tremendous during the last decade. Along with this trend, the community has been searching, investigating and looking for programming models that provide both control on the data locality and flexibility of the data handling.

SHMEM was introduced by Cray [8] in 1994, followed shortly later by SGI [19]. In an effort to provide a homogeneous, portable standard for the language, the OpenSHMEM consortium released a specification for the application programming interface [16]. The final version of OpenSHMEM 1.0 was released in January 2012.

The OpenSHMEM standard is a programming paradigm for parallel applications that uses single-sided communications. It opens gates for exciting research in distributed computing on this particular communication model.

This paper presents Paris OpenSHMEM (POSH), which is a portable, open-source implementation of OpenSHMEM. It uses a high-performance communication engine on shared memory based on the Boost library [10], and benefits from the template engine provided by current C++ compilers.

This paper describes the implementation choices and the algorithms that have been used in POSH in order to fit with the memory model and the communication model while obtaining good performance and maintaining portability.

This report is organized as follows; section 2 gives a short overview of the related literature about parallel programming paradigms and distributed algorithms on shared memory and one-sided communication models. Section 3 gives details about the memory model and the communication model which are considered here. Section 4 describes the implementation choices that were made in POSH. Section 5 presents some performance results that were obtained by the current version of POSH. Last, section 6 concludes the report and states some open issues and

*Some experiments presented in this paper were carried out using the Grid'5000 experimental testbed, being developed under the INRIA ALADDIN development action with support from CNRS, RENATER and several Universities as well as other funding bodies (see https://www.grid5000.fr).

future works that will be conducted on POSH.

2 Related works

Traditionally, distributed systems are divided into two categories of models for their communications: those that communicate by sending and receiving messages (*i.e.,* message-passing systems) and those that communicate using registers of shared memory where messages are written and read from (*i.e.,* shared memory systems) [21, 11].

Along with the massive adoption of the many-core hardware architecture, researchers and engineers have tried to find the most efficient programming paradigm for such systems. The idea is to take advantage of the fact that processing units (processes or threads) have access to a common memory: those are *shared memory* systems. Unix IPC V5 and posix threads are the most basic programming tools for that. OpenMP [9] provides an easy-to-use programming interface and lets the programmer write programs that look very similar to sequential ones, and the parallelization is made by the compiler. Therefore, the compiler is in charge with data decomposition and accesses. Several data locality policies have been implemented to try to make the best possible guess about where it must be put to be as efficient as possible [22]. OpenMP performs well on regular patterns, where the data locality can be guessed quite accurately by the compiler. Cilk [2] and TBB [18] can also be cited as programming techniques for shared-memory systems.

MPI [12, 14] has imposed itself as the *de facto* programming standard for distributed-memory parallel systems. It is highly portable, and implementations are available for a broad range of platforms. MPICH [15] and Open MPI [13] must be cited among the most widely used open-source implementations. It can be used on top of most local-area communication networks, and of course most MPI implementations provide an implementation on top of shared memory. MPI is often referred to as "the assembly language of parallel computing": the programmer has total control of the data locality, however all the data management must be implemented by hand by the programmer. Moreover, it is highly synchronous: even though specific one-sided communications have been introducted in the MPI2 standard [14], the sender and the receiver must be in matching communication routines for a communication to be performed.

Hence, there exists two opposing trends in parallel programming techniques: programming easiness versus data locality mastering. A third direction exists and is becoming pertinent with many-core architectures: making data locality mastering easier and more flexible for the programmer. UPC can be cited as an example of programming technique that is part of that third category [7]. It provides compiler-assisted automatic loops, automatic data repartition in (potentially distributed) shared memory, and a set of one-sided communications.

One-sided communications are natural on shared memory systems, and more flexible than two-sided communications in a sense that they do not require that both of the processes involved in the communication (origin and destination of the data) must be in matching communication routines. However, they require a careful programming technique to maintain the consistency of the shared memory and avoid race conditions [6].

SHMEM was introduced by Cray [8] as part of its programming toolsuite with the Cray T3 series, and SGI created its own dialecte of SHMEM [19].

Some implementations also exist for high-performance RDMA networks: Portals has been working on a specific support for OpenSHMEM by their communication library [1]. Some other implementations are built on top of MPI implementations over RDMA networks, such as [4] for Quadrics networks or [17] over InfiniBand networks.

In this paper, we propose to use a shared memory communication engine based on the Boost.Interprocess library [10], which is itself using the POSIX `shm` API.

3 Memory model

OpenSHMEM considers a memory model where every process of the parallel application owns a local bank of memory which is split into two parts:

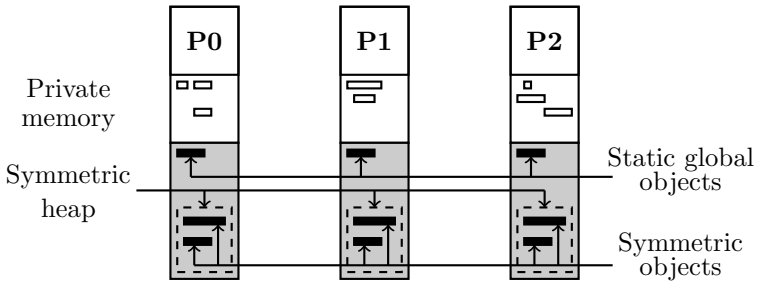

Figure 1: Memory organization with global static objects and data in the symmetric heap. Static objects are remotely accessible, dynamic objects are located in the symmetric heap.

- Its *private memory*, which is accessible by itself only; no other process can access this area of memory.
- Its *public memory*, that can be accessed by any process of the parallel application, in read/write mode.

This memory organization is represented in figure 1. Each process owns its private memory (white rectangles) as well as an area of public memory (gray rectangles).

3.1 Symmetric objects

The public memory of each process is defined as a *symmetric heap*. This notion of symmetry is important because it is a necessary condition for some helpful memory-management properties in OpenSHMEM (see section 4.1.2 and 4.5.3). It means that for any object which is stored in the symmetric heap of a process, there exists an object of the same type and size and the same address, in the symmetric heap of all the other processes of the parallel application.

Dynamically-allocated variables, *i.e.,* variables that are allocated explicitely at run-time, are placed in the symmetric heap. OpenSHMEM provides some functions that allocate and deallocate memory space dynamically in the symmetric heap.

Another kind of data is put into processes' public memory: global, static variables (in C/C++).

All the data that is placed in the processes' symmetric heap and all the global, static variables are remotely accessible by all the other processes. These two kinds of variables are represented in figure 1. The gray areas prepresent the public memory of each process; composed with global static variables and the symmetric heap. The white areas represent the private memory of each process.

3.2 One-sided communications

Point-to-point communications in OpenSHMEM are *one-sided*: a process can reach the memory of another process without the latter knowing it. It is very convenient at first glance, because no synchronization between the two processes is necessary like with two-sided communications. However, such programs must be programmed very carefully in order to maintain memory consistency and avoid potential bugs such as race conditions.

Point-to-point communications in OpenSHMEM are based on two primitives: `put` and `get`.

- A `put` operation consists in writing some data at a specific address of remote process's public memory.
- A `get` operation consists in reading some data, or fetching it, from a specific address of a remote process's public memory.

Data movements are made between the public memory of the local process and the private memory of the remote process: a process reads a remote value and stores it in its own private memory, and it writes the value of a variable located in its own private memory into the public memory of a remote pro-

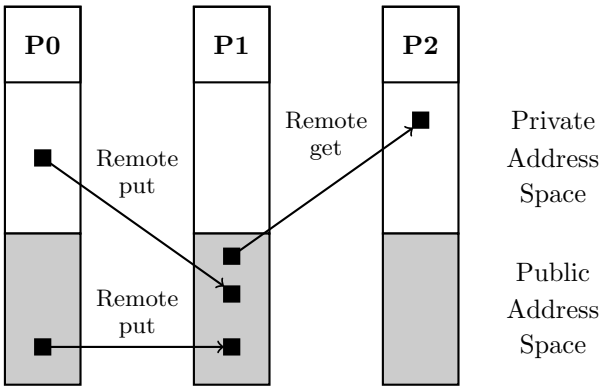

Figure 2: Put and get communications between three processes.

cess. That memory model and one-sided communications performed on this model have been described more thoroughly in [5, 6].

These operations between the public and private memory areas of three processes of a parallel application are represented in figure 2.

4 Implementation details

In this section, we describe the implementation details and the design choices that have been made in POSH.

4.1 Shared memory communication engine

POSH relies on Boost's library for inter-process communications `Boost.Interprocess`. In particular, it is using the `managed_shared_memory` class.

Basically, each process's shared heap is an instance of `managed_shared_memory`. Data is put into that heap by an allocation method provided by this class followed by a memory copy. Locks and all the manipluation functions related to the shared memory segment are also provided by Boost.

4.1.1 Memory management

Allocation and deallocation Memory can be allocated in the symmetric heap with the `shmalloc` function. Internally, that function calls the `allocate` function of the `managed_shared_memory` class on the process's shared heap. That class also provides an `allocate_aligned` routine which is called by the `shmemalign` routine. Memory is freed by `shfree` by using a call to the function `deallocate` provided by `managed_shared_memory`.

Remote vs local address of the data These three SHMEM functions are defined as *symmetric* functions: all the processes must call them at the same time. They are required by the OpenSHMEM standard to perform a global synchronization barrier before returning. As a consequence, if all the memory allocations and deallocations have been made in a symmetric way, a given chunk of data will have the same local address in the memory of all the processes. Hence, we can access a chunk of data on process A using the address it has on process B. That property is extremely useful for remote data accesses.

4.1.2 Access to another process's shared heap

Fact 1. *If all the processing elements are running on the same architecture, the offset between the beginning of a symmetric heap and a symmetric object which is contained by this heap is the same on each processing element.*

Proof. Thanks to the property of symmetry between

the shared heaps, data is put into the symmetric heaps at the same *local* address across processes. For instance, if a variable is located in the shared heap of a process at address X, the same variable (possibly with a different value) will be located at the same address X in the symmetric heap of all the other processes.

That property is enforced by the fact that, as stated in the OpenSHMEM standard, memory allocations which are performed in the symmetric heaps end by a call to a global synchronization barrier. As a consequence, all the processing elements *must* allocate space. If they do not allocate the same space, then this is a mistake made by the programmer and the behavior of the program is undefined (paragraph 6.4 of the OpenSHMEM standard). □

Corollary 1. *As a consequence of fact 1, each processing element can compute the address of a variable located in another processing element's symmetric heap by using the following formula:*

$$addr_{remote} = heap_{remote} + (addr_{local} - heap_{local}) \quad (1)$$

Where $addr_{remote}$ and $addr_{local}$ respectively denote the address of symmetric objects that are located in the symmetric heap of a remote processing element and in the local symmetric heap, and $heap_{remote}$ and $heap_{local}$ respectively denote the address at which the other processing element's heap and the local heap are mapped in the local memory.

When a given process wants to access another process's symmetric heap, it proceeds as follows:

- Build the remote symmetric heap's name, based on its rank;
- Make sure the remote symmetric heap exists. If it does not exist yet, we wait a little bit and try again;
- Build a local object that contains a reference to the shared memory and maps it into the local memory;
- Then find out what the address of the remote object is in the local memory. We use a little trick here. Boost provides the notion of *handle* to locate a chunk of memory in a shared memory segment. This handle is relative to a shared memory segment. So we get the handle to the chunk of data which is in *our own* symmetric heap, and then we get the address that corresponds to this handle in the *remote* symmetric heap. This trick is allowed by the property of symmetry between all the symmetric heaps, which is explained in section 4.1.1.

That being done, we have direct access to the remote symmetric heap. In particular, we can copy data from and to this remote symmetric heap.

However, building the remote heap's name and the corresponding shared object is quite expensive in terms of object crations (and destructions at the end of this process). As a consequence, they are all created at startup-time and cached in a local structure (a table). This operation in itself is quite inexpensive thanks to move mechanisms that are featured by C++11. Hence, when a process needs to access another process's heap, it simply looks into this local table for the reference to this segment of shared memory and accesses it.

4.2 Symmetric static data

The memory model specifies that global, static variables are made accessible for other processes. In practice, these variables are placed in the BSS segment if they are not initialized at compile-time and in the data segment if they are initialized. Unfortunately, there is no simple way to make these areas of memory accessible for other processes.

Therefore, POSH uses a small trick: we put them into the symmetric heap at the very beginning of the execution of the program, before anything else is done.

A specific program, called the *pre-parser*, parses the source code and searches for global variables that are declared as static. It finds out how they must be allocated (size, etc) and generates the appropriate allocation/deallocation code lines.

When the OpenSHMEM library is initialized (*i.e.,* when the `start_pes` routine is called), it dumps the allocation code into the source code. When the program exits (*i.e.,* when the keyword `return` is found

in the main function), the deallocation code lines are inserted before each `return` keyword.

4.3 Datatype-specific routines

OpenSHMEM defines a function for each data type. For example, fetching a single variable can be done by:

- `short shmem_short_g( short *addr, int pe )` for variables of type `short`
- `int shmem_int_g( int *addr, int pe )` for variables of type `int`
- `long shmem_long_g( long *addr, int pe )` for variables of type `long`
- `float shmem_float_g( float *addr, int pe )` for variables of type `float`
- `double shmem_double_g( double *addr, int pe )` for variables of type `double`
- `long long shmem_longlong_g( long long *addr, int pe )` for variables of type `long long`
- `long double shmem_longdouble_g( long double *addr, int pe )` for variables of type `long double`

A large part of this code can be factorized by using an extremely powerful feature of the C++ language: templates. The corresponding code is written only once, and then the template engine instanciates one function for each data type. Hence, only one function needs to be written.

In the aforementioned example, only one function was written:

```
template<class T> T shmem_template_g(
    T* addr, int pe );
```

That function is called by each of the OpenSHMEM `shmem_*_...` functions. Each call is actually a call to the compiler-generated function that uses the adequate data type. That function is generated at compile-time, not at run-time: consequently, calling that function is just as fast as if it had been written manually.

4.4 Peer-to-peer communications

Peer-to-peer communications are using memory copies between local and shared buffers. As a consequence, memory copy is a highly critical matter of POSH. Several implementations of `memcpy` are featured by POSH in order to make use of low-level hardware capabilities such as MMX, MMX2, SSE or SSE2 instruction sets, or the default `memcpy` provided by the kernel.

One of these implementations is activated by using a compiler directive. In order to minimize the number of conditional branches, selecting one particular implementation is made at compile-time rather than at run-time.

A comparison of the performance obtained by these different implementations is presented in section 5.

4.5 Collective communications

Collective communications rely on point-to-point communications that perform the actual inter-process data movements. Two options are available for these point-to-point communications:

- *Put-based* communications push the data into the next processes;
- *Get-based* communications pull the data from other processes.

4.5.1 Data structure for collective communication

Each process holds a data structure in their shared heap (hence, other processes can access it). This data structure contains information about the ongoing collective operation:

- A pointer to the buffer that contains the data that is moved by the collective operation;
- A counter, that counts how many remote processes have accessed the local data;
- A type, that keeps what collective operation is underway;
- A boolean that specifies whether or not the collective communication is already in progress;

- In debug and in safe mode we can keep the size of the data buffer, in order to check that the allocated buffer has the same size as the data we are trying to put into it.

This data structure is intialized during the initalization of the OpenSHMEM library, after the initalization of the symmetric heap. It is reset at the end of each collective communication, in order to make sure the place is "clean" for the next collective communication.

4.5.2 Progress of a collective operation

The communication model used by OpenSHMEM and its point-to-point communication engine is particular in a sense that it is using *one-sided* operations. As a consequence, a process can access in read or write mode another process's memory without the knowledge of the latter process. One consequence of this fact is that a process can be involved in a collective communication without having actually entered the call to the corresponding routine yet.

Hence, if a process A must access the symmetric heap of a process B, the former process must check whether or not the latter has entered the collective communication yet. A boolean variable is included in the collective data structure for this purpose.

If the remote process has not entered the collective communication yet, its collective data structure must be initialized remotely.

If some data must be put into a remote process that has yet to initialize its collective data structure, we only copy the pointer to the shared source buffer. The actual memory allocation will be made later. However, only temporary memory allocations are made within collective operations. Buffers that are used as parameters of a collective operation (source, target and work arrays) must be allocated before the call to this operation. Since memory allocations end by a global barrier, no processing element can enter a collective operation if not all of them have finished their symmetric memory allocations.

When a process enters a collective operation, it checks whether the operation is already underway, *i.e.,* whether its collective data structure has already been modified by a remote process. If so, we need to make the actual memory allocation for the local data and copy what has already been put somewhere in another shared memory area.

A process exits the collective communication as soon as its participation to the communication is over. Hence, no other process will access its collective data structure. It can therefore be reset.

4.5.3 Temporary allocations in the shared heap

With some collective communication algorithms, it can be necessary to allocate some temporary space in the shared heap of a given processing element. However, if we allocate some memory in one heap only, we break the important symmetry assumption made in section 4.1.1. Nevertheless, we will see here that actually, they have no impact in the symmetry of the shared heaps outside of the affected collective operation.

Lemma 1. *Non-symmetric, temporary memory allocations in the heap of a subset of the processing elements that are performed during collective operations do not break the symmetry of the heaps outside of the concerned collective operation.*

Proof. Semantically, collective operations are symmetric, in a sense that all the concerned processing elements must take part of them. As a consequence, if all the heaps are symmetric before they enter the collective operation and if there is no difference between the state of each heap at the beginning and at the end of the collective operation, hence, the symmetry is not broken.

Buffers (source, target and work arrays) must be allocated by a call to `shmalloc` before calling the collective operation. Since `shmalloc` ends by a call to a global barrier, no processing element can enter the collective operation before all the processing elements have performed the last memory allocation before this collective call.

According to the sematics of collective operations, at any given moment of the progress of a collective operation, the relative state of each pair of processes A and B is one of the following:

- Neither A nor B have entered the collective operation yet, or both of them have left it: in this case, A and B do not interact with each other's symmetric heaps, and memory allocations are performed by symmetric allocations, as described in section 4.1.1.
- Both processes A and B are inside the call to the collective routine and can interfere with each other's symmetric heap: anything can happen in the symmetric heaps within the progress of the collective operation, as long as non-symmetric operations are cleaned-up when each process leaves the collective operation.
- A process A can push some data into the memory of another process B that has not entered the collective operation yet. In this case, B is unknowingly taking part of the collective operation, but its shared heap can potentially be modified by other processes. However, as stated above, B cannot perform any symmetric memory allocation between that moment and the moment when it enters the collective operation. As a consequence, the non-symmetric allocations that can potentially be made in B's shared heap by other processing elements while the latter are taking part of the collective operation do not break the symmetry as long as they are reverted (*i.e.,* freed) before B leaves the collective operation.
- A process A still part of the collective operation whereas another process B is done with its participation to the operation and therefore, B has left the call to the collective routine, which means that its participation to the collective operation is over. As a consequence, other processes have no reason to modify its shared heap.

□

4.5.4 Switching between algorithms

In order to reduce the number of conditional branches, collective communication algorithms are chosen at compile-time. The choice is made using compiler variables and conditions. A default choice is provided if no option is passed to the compiler, and a warning is displayed.

4.5.5 Run-time error checking

Collective communications are an important source of bugs and errors in parallel programs. When compiled in safe mode, the OpenSHMEM library provides some run-time error checking facilities.

For instance, it can check whether the size of the available buffer is equal to the size of the data that is about to be pushed into or pulled from a shared collective data structure. It can also make sure that the collective data structures of the local and the remote processes are performing the same type of collective operation.

4.6 Locks and atomic operations

Boost provides *named mutexes* for locking purpose. These locks are interesting, because they can be specific for a given symmetric heap. Each process uses the same given name for a given chunk of data on a given symmetric heap. Using a mutex that locally has the same name as all the other local mutexes, processes ensure mutual exclusion. Hence, we can make sure that a chunk of data is accessed by one process only.

Boost also provides a function that executes a given function object *atomically* on a managed shared memory segment such as the one that is used here to implement the symmetric heap. Hence, we can perform atomic operations on a (potentially remote) symmetric heap.

4.7 Run-time environment

As with any parallel program, the run-time environment of OpenSHMEM is here to:

- Spawn the parallel processes;
- Make sure they know how to communicate with each other;
- Monitor them, and take the appropriate actions if one of them dies;
- Terminate the execution when necessary;
- Forward the IOs and signals through the gateway process that provides the user with an access to the parallel execution.

Process spawning Processes are spawned individually by separate threads. At first, a pool of threads is created: the `workers` thread group. Then each thread forks a process: the corresponding OpenSHMEM processing element. The master thread then yields its slice of time (`sched_yield`) and waits on a condition. Eventually, the threads are joined.

Contact information Processes communicate with each other using shared memory segments, which are their symmetric shared heap. The name of this shared memory segment is built using a constant basis and the rank of the target process. Hence, processes can communicate with each other as soon as they know their rank.

Inputs and outputs The run-time environment is supposed to forward IOs and signals between the user and the parallel application. More specifically, the parallel application is made of several processes, whereas the user is in contact with only one process: the master process, which is used as a gateway between the user and the application.

For instance, if a parallel process performs an output (`printf`, `std<<cout`...), the result of this output will be displayed to the user by the gateway process. Similarly, if the user sends a signal to the gateway process (*e.g.,* `SIGKILL`), this signal is sent to all the processes of the parallel application.

The mechanism used here to create the parallel processes preserves IOs. The parallel processes are offsprings of the gateway processes: hence, their IOs are forwarded by default.

If necessary, `stdout`, `stderr` and `stdin` can be duplicated and copied just before the `execve` system call.

Run-time debugging Parallel processes can require to be debugged in an interactive way at runtime. In this case, a sequential debugger like `gdb` can be attached to a given process.

To allow this attachment, the parallel process is stuck in an infinite loop at the begining of its initialization.

Debugging information and checks have to be placed in parts of code that are removed by the compiler when the debugging mode is disabled. The compiler variable is called `_DEBUG`.

POSH also provides a *safe* mode. This mode enables some debugging and error checking information for the parallel program, whereas the debug mode enables debugging information for the OpenSHMEM library. The compiler variable is called `_SAFE`.

For instance, the safe mode checks that when a process wants to run a collective communication, it is not already participating to another collective communication.

In order to be able to choose whether to use it or not without affecting the performance, we chose not to make it a run-time option, but a compile-time option. As a consequence, it must be enabled by a compiler option when the OpenSHMEM library is compiled.

Code related to the safe mode is left out by the compiler when it is not meant to be enabled. We are using a compiler variable called `_SAFE`.

5 Performance and experimental results

This section presents some evaluations of the performance achieved by POSH. Time measurements were done using `clock_gettime()` on the `CLOCK_REALTIME` to achieve nanosecond precision. All the programs were compiled using `-Ofast` if available, `-O3` otherwise. Each experiment was repeated 20 times after a warm-up round. We measured the time taken by data movements (put and get operations) for various buffer sizes.

5.1 Memory copy

Since memory copy (`memcpy`-like) is a highly critical function of POSH, we have implemented several versions of this routine and evaluated them in a separate micro-benchmark. The compared performance of these various implementations of `memcpy()` is out

	Memory copy latency (ns)			
	memcpy	MMX	MMX2	SSE
Caire	38.85	41.10	38.65	**38.05**
Jaune	1277.90	1273.90	**1269.90**	1279.90
Magi10	45.40	**38.20**	39.90	40.70
Maximum	21.70	**20.25**	20.45	21.00
Pastel	**1997.30**	1997.40	2011.35	1997.35

	Memory copy bandwidth (Gb/s)			
	memcpy	MMX	MMX2	SSE
Caire	**18.40**	12.25	18.18	18.37
Jaune	9.84	10.03	16.44	**16.60**
Magi10	**22.93**	21.13	17.06	20.77
Maximum	67.47	47.52	76.59	**77.91**
Pastel	20.27	9.12	**20.32**	19.82

Table 1: Comparison of the performance observed with various memcpy implementations

of the scope of this paper. A good description of this challenge, the challenges that are faced and the behavior of the various possibilities can be found in [20].

The goal of POSH is to achieve high-performance while being portable. As a consequence, the choice has been made to provide different implementations of `memcpy` and let the user choose one at compile-time, while providing a default one that achieves reasonably good performance across platforms.

We have compared several implementations of `memcpy` on various platforms that feature different CPUs: an Intel Core i7-2600 CPU running at 3.40GHz (Maximum), a Pentium Dual-Core CPU E5300 running at 2.60GHz (Caire), an AMD Athlon 64 X2 Dual Core Processor 5200+ (Jaune), a large NUMA node featuring 4 CPUs with 10 physical cores each (20 logical cores with hyperthreading), Intel Xeon CPU E7-4850 running at 2.00GHz (Magi10) and a NUMA node featuring 2 CPUs with 2 cores each, Dual-Core AMD Opteron Processor 2218 running at 2.60GHz (Pastel). All the platforms are running Linux 3.2, except Jaune (2.6.32) and Maximum (3.9). The code was compiled by gcc 4.8.2 on Maximum, Caire and Jaune, gcc 4.7.2 on Pastel and icc 13.1.2 on Magi10.

We compared the stock `memcpy()` and MMX-, MMX2- and SSE-based implementations. The performance (latency and bandwidth) of these implementations on the aforementioned platforms are summarized in table 1.

We can see that the variations between latencies obtained by all the four implementations are very small, except for Magi10 (the large NUMA node). Pastel (the Opteron node) features a slightly better bandwidth with MM2, whereas Jaune and Maximum (the Athlon XP and the Core i7 nodes) achieve higher bandwidth with SSE and Caire and Magi10 (the Dual-Core and the large NUMA node) perform better with the stock `memcpy`. The large performance gap on Jaune may be explained by the relatively old software it is running. Overall, the stock `memcpy` performs quite well (best performance of close to the best performance), except on Jaune and Maximum, on which the bandwidth is largely improved by using SSE instructions.

5.2 POSH communication performance

We evaluated the communication performance obtained with POSH. On Caire, Magi10 and Pastel, we used the stock `memcpy` for data movements. On Jaune and Maximum, we used both the SSE-based implementation and the sock `memcpy`. Table 2 present the latency and bandwidth obtained by put and get operations with POSH. Figure 3 plots the latency and bandwidth obtained on Maximum.

On the "fast" machines (Caire, Magi10 and Maximum), the latency is too small to be measured precisely by our measurement method. We can see on table 2 that the latency has the same order of magnitude as the one obtained by a `memcpy` within the memory space of a single process. However, measuring the latency on regular communication patterns gives an indication on the overall latency of the communication engine, but may be different from what would be obtained on more irregular patterns, where the segment of shared memory is not in the same

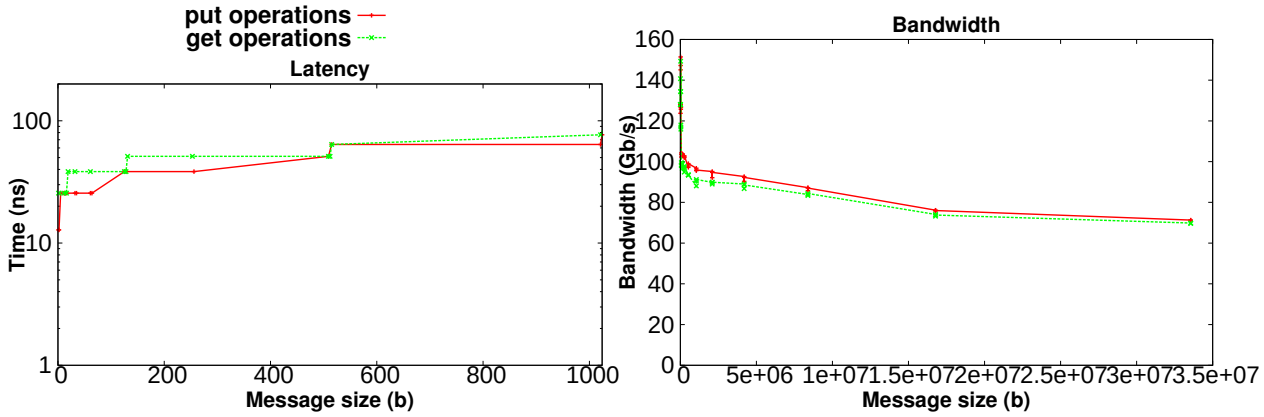

Figure 3: Communication performance of Paris OpenSHMEM on Maximum.

	SHMEM latency (ns)			
	Best copy		`memcpy`	
	get	put	get	put
Caire	38.40	38.40	38.40	38.40
Jaune	1741.85	1665.90	1667.90	1663.90
Magi10	38.40	38.40	38.40	38.40
Maximum	38.40	38.40	38.40	38.40
Pastel	1830.40	1689.60	1830.40	1689.60

	SHMEM bandwidth (Gb/s)			
	Best copy		`memcpy`	
	get	put	get	put
Caire	18.36	18.38	18.36	18.38
Jaune	17.62	17.55	10.52	10.59
Magi10	20.46	20.16	20.46	20.16
Maximum	74.09	76.15	68.51	69.28
Pastel	26.07	25.50	26.07	25.50

Table 2: Comparison of the performance observed by put and get operations with POSH

cache memory as the process that performs the data movement. In the latter case, the kernel's scheduling performance is highly critical.

Similarly, the bandwidth obtained by POSH has little overhead compared with the one obtained by a `memcpy` within the memory space of a single process. We can conclude here that our peer-to-peer communication engine adds little overhead, no to say a negligible one, and inter-process communications are almost as fast as local memory copy operations.

5.3 Comparison with another communication library

We used a similar benchmark to evaluate the communication performance of Berkeley UPC, whose communication engine, GASNet [3], uses `memcpy` to move data. As a consequence, the results obtained here must be compared to those obtained in the previous sections with the stock `memcpy`. Here again, we can see that BUPC inter-process data movement operations have little overhead compared to a memory copy that would be performed within the memory space of a single process. The results are presented in table 3.

We can see here that both POSH and another one-sided communication library (Berkeley UPC) have performance that are close to a memory copy within the address space of a single process. Besides, we have seen how the performance can benefit from a tuned memory copy routine.

	UPC latency (ns)	
	get	put
Caire	39.40	37.55
Jaune	1623.90	1623.90
Magi10	73.80	54.90
Maximum	26.75	25.00
Pastel	2025.10	1689.95

	UPC bandwidth (Gb/s)	
	get	put
Caire	18.03	18.45
Jaune	9.95	10.63
Magi10	18.64	16.33
Maximum	67.45	68.86
Pastel	23.52	25.06

Table 3: Comparison of the performance observed by put and get operations with UPC

6 Conclusion and perspective

In this paper, we have presented the design and implementation of POSH, an OpenSHMEM implementation based on a shared memory engine provided by Boost.Interprocess, which is itself based on the POSIX `shm` API. We have presented its architecture, a model for its communications and proved some properties that some implementation choices rely on. We have presented an evaluation of its performance and compared it with a state-of-the-art implementation of UPC, another programming API that follows the same communication model (one-sided communications).

We have seen that POSH achieves a performance which is comparable with this other library and with simple memory copies. We have also shown how it can be tuned in order to benefit from optimized low-level routines.

That communication model opens perspectives on novel work on distributed algorithms. The architectural choices that were made in POSH make it possible to use it as a platform for implementing and evaluating them in practice. For instance, locks, atomic operations and collective operations are classical problems in distributed algorithms. They can be reviewed and re-examined in that model in order to create novel algorithms with an original approach.

References

[1] Brian W. Barrett, Ron Brightwell, K. Scott Hemmert, Kevin T. Pedretti, Kyle B. Wheeler, and Keith D. Underwood. Enhanced support for OpenSHMEM communication in Portals. In *Hot Interconnects*, pages 61–69, 2011.

[2] Robert D Blumofe, Christopher F Joerg, Bradley C Kuszmaul, Charles E Leiserson, Keith H Randall, and Yuli Zhou. *Cilk: An efficient multithreaded runtime system*, volume 30. ACM, 1995.

[3] Dan Bonachea. GASNet Specification, v1.1. Technical Report UCB/CSD-02-1207, U.C. Berkeley, 2002.

[4] Ron Brightwell. A new MPI implementation for Cray SHMEM. In *Recent Advances in Parallel Virtual Machine and Message Passing Interface, 11th European PVM/MPI Users' Group Meeting (EuroPVM/MPI'04)*, pages 122–130, 2004.

[5] Franck Butelle and Camille Coti. A model for coherent distributed memory for race condition detection. In *proceedings of the 13th Workshop on Advances in Parallel and Distributed Computational Models (APDCM'11)*, pages 579–585, Anchorage, Ak, May 2011.

[6] Frank Butelle and Camille Coti. Data coherency in distributed shared memory. *IJNC*, 2(1):117–130, 2012.

[7] UPC Consortium. UPC Language Specifications, v1.2. Technical Report LBNL-59208, Lawrence Berkeley National, 2005.

[8] Inc. Cray Research. SHMEM Technical Note for C. Technical Report SG-2516 2.3, 1994.

[9] Leonardo Dagum and Ramesh Menon. OpenMP: an industry standard API for

shared-memory programming. *IEEE Comput. Sci. Eng.*, 5:46–55, January 1998.

[10] Boost Documentation. Chapter 9. boost interprocess. Available at `http://www.boost.org/doc/libs/1_55_0/doc/html/interprocess.html`.

[11] S. Dolev. *Self-Stabilization*. MIT Press, 2000.

[12] Message Passing Interface Forum. MPI: A message-passing interface standard. Technical Report UT-CS-94-230, Department of Computer Science, University of Tennessee, April 1994. Tue, 22 May 101 17:44:55 GMT.

[13] Edgar Gabriel, Graham E. Fagg, George Bosilca, Thara Angskun, Jack J. Dongarra, Jeffrey M. Squyres, Vishal Sahay, Prabhanjan Kambadur, Brian Barrett, Andrew Lumsdaine, Ralph H. Castain, David J. Daniel, Richard L. Graham, and Timothy S. Woodall. Open MPI: Goals, concept, and design of a next generation MPI implementation. In *Recent Advances in Parallel Virtual Machine and Message Passing Interface, 11th European PVM/MPI Users' Group Meeting (EuroPVM/MPI'04)*, pages 97–104, Budapest, Hungary, September 2004.

[14] Al Geist, William D. Gropp, Steven Huss-Lederman, Andrew Lumsdaine, Ewing L. Lusk, William Saphir, Anthony Skjellum, and Marc Snir. MPI-2: Extending the message-passing interface. In Luc Bougé, Pierre Fraigniaud, Anne Mignotte, and Yves Robert, editors, *1st European Conference on Parallel and Distributed Computing (EuroPar'96)*, volume 1123 of *Lecture Notes in Computer Science*, pages 128–135. Springer, 1996.

[15] William Gropp. MPICH2: A New Start for MPI Implementations. In Dieter Kranzlmller, Jens Volkert, Peter Kacsuk, and Jack Dongarra, editors, *Recent Advances in Parallel Virtual Machine and Message Passing Interface*, volume 2474 of *Lecture Notes in Computer Science*, pages 37–42. Springer Berlin / Heidelberg, 2002.

[16] High Performance Computing Tools group at the University of Houston and Oak Ridge National Laboratory Extreme Scale Systems Center. OpenSHMEM application programming interface, version 1.0 final. `http://www.openshmem.org`, January 2012.

[17] Jiuxing Liu, Jiesheng Wu, and D. K. Panda. High Performance RDMA-Based MPI Implementation over InfiniBand. In *17th Annual ACM International Conference on Supercomputing (ICS'03)*, 2003.

[18] James Reinders. *Intel threading building blocks: outfitting C++ for multi-core processor parallelism*. O'Reilly Media, Inc., 2010.

[19] SGI. SHMEM API for Parallel Programming . `http://www.shmem.org`.

[20] Zack Smith. Bandwidth: a memory bandwidth benchmark for x86 / x86_64 based Linux/Windows/MacOSX. `http://zsmith.co/bandwidth.html`, June 2010.

[21] Gerard Tel. *Introduction to Distributed Algorithms*. Cambridge University Press, 1994.

[22] Tien-Hsiung Weng and Barbara M Chapman. Implementing OpenMP Using Dataflow Execution Model for Data Locality and Efficient Parallel Execution. In *Proceedings of the 16th International Parallel and Distributed Processing Symposium*, IPDPS '02, pages 180–, Washington, DC, USA, 2002. IEEE Computer Society.